BITS

computing & communications news

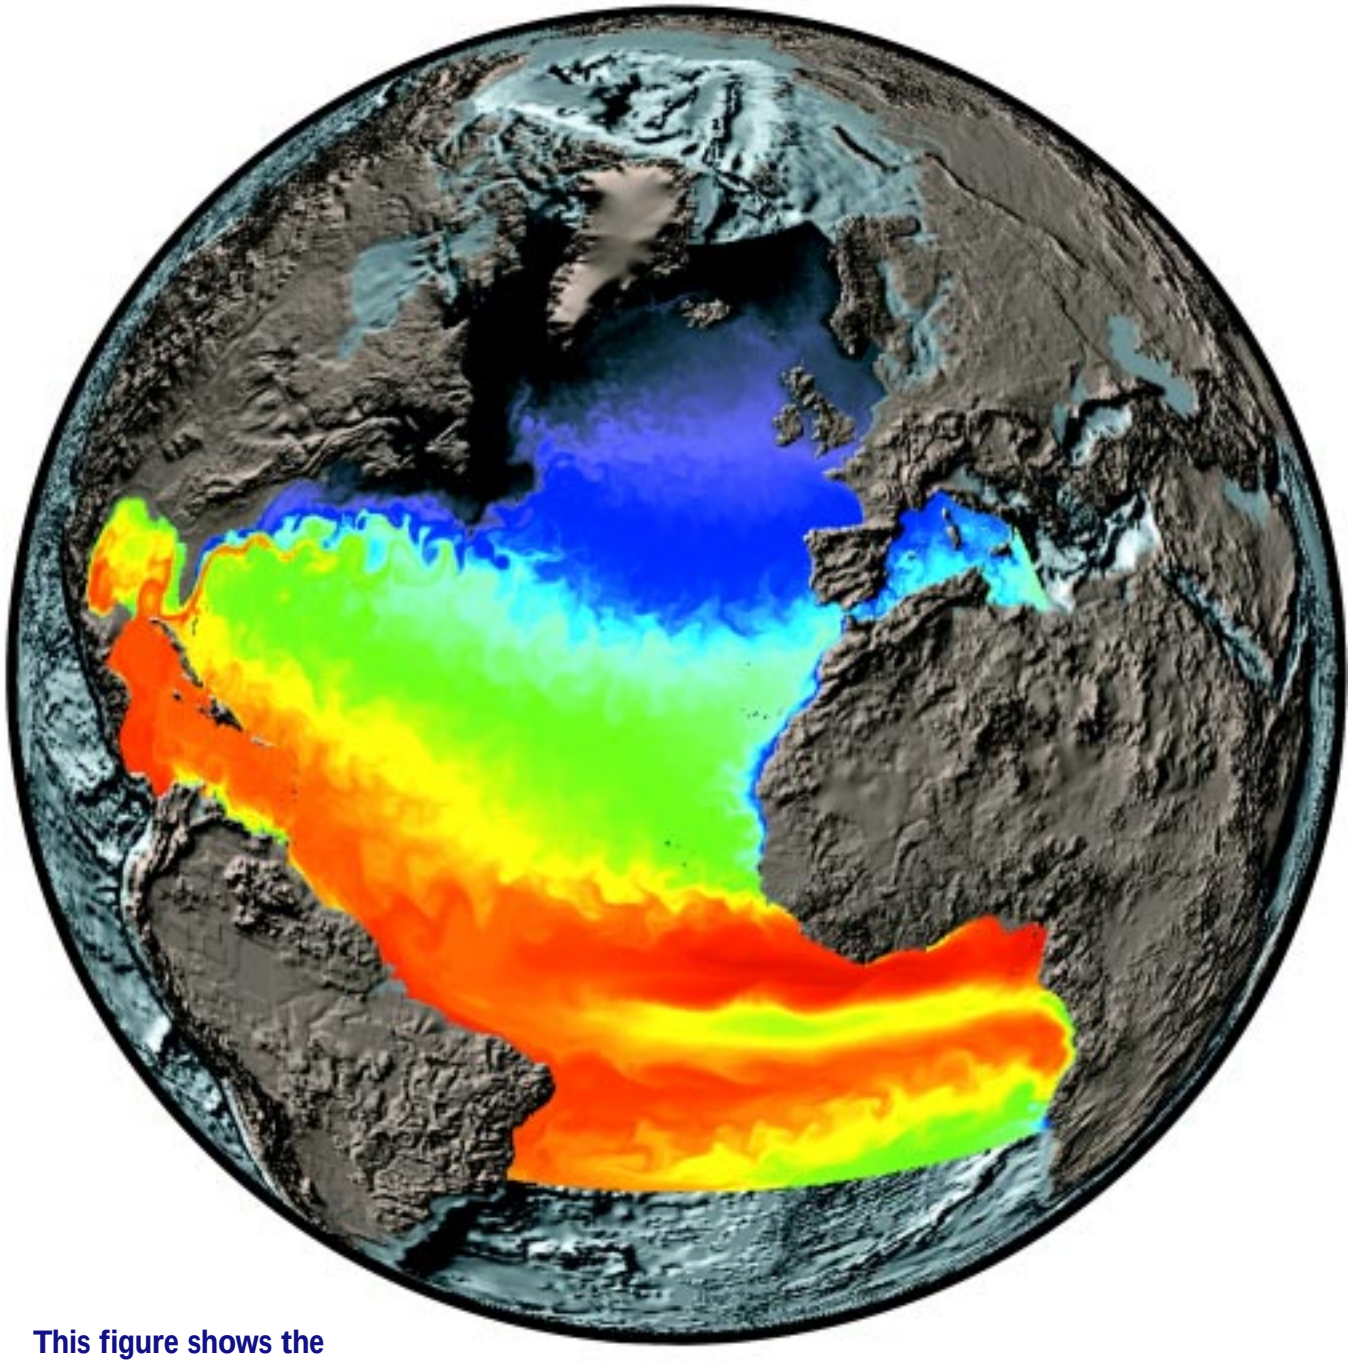

This figure shows the
sea surface temperature from a
0.1 degree simulation of the North Atlantic
Ocean computed by the Parallel Ocean Program
and visualized with the POPTEX post-processing software.

Los Alamos
NATIONAL LABORATORY

January 2000

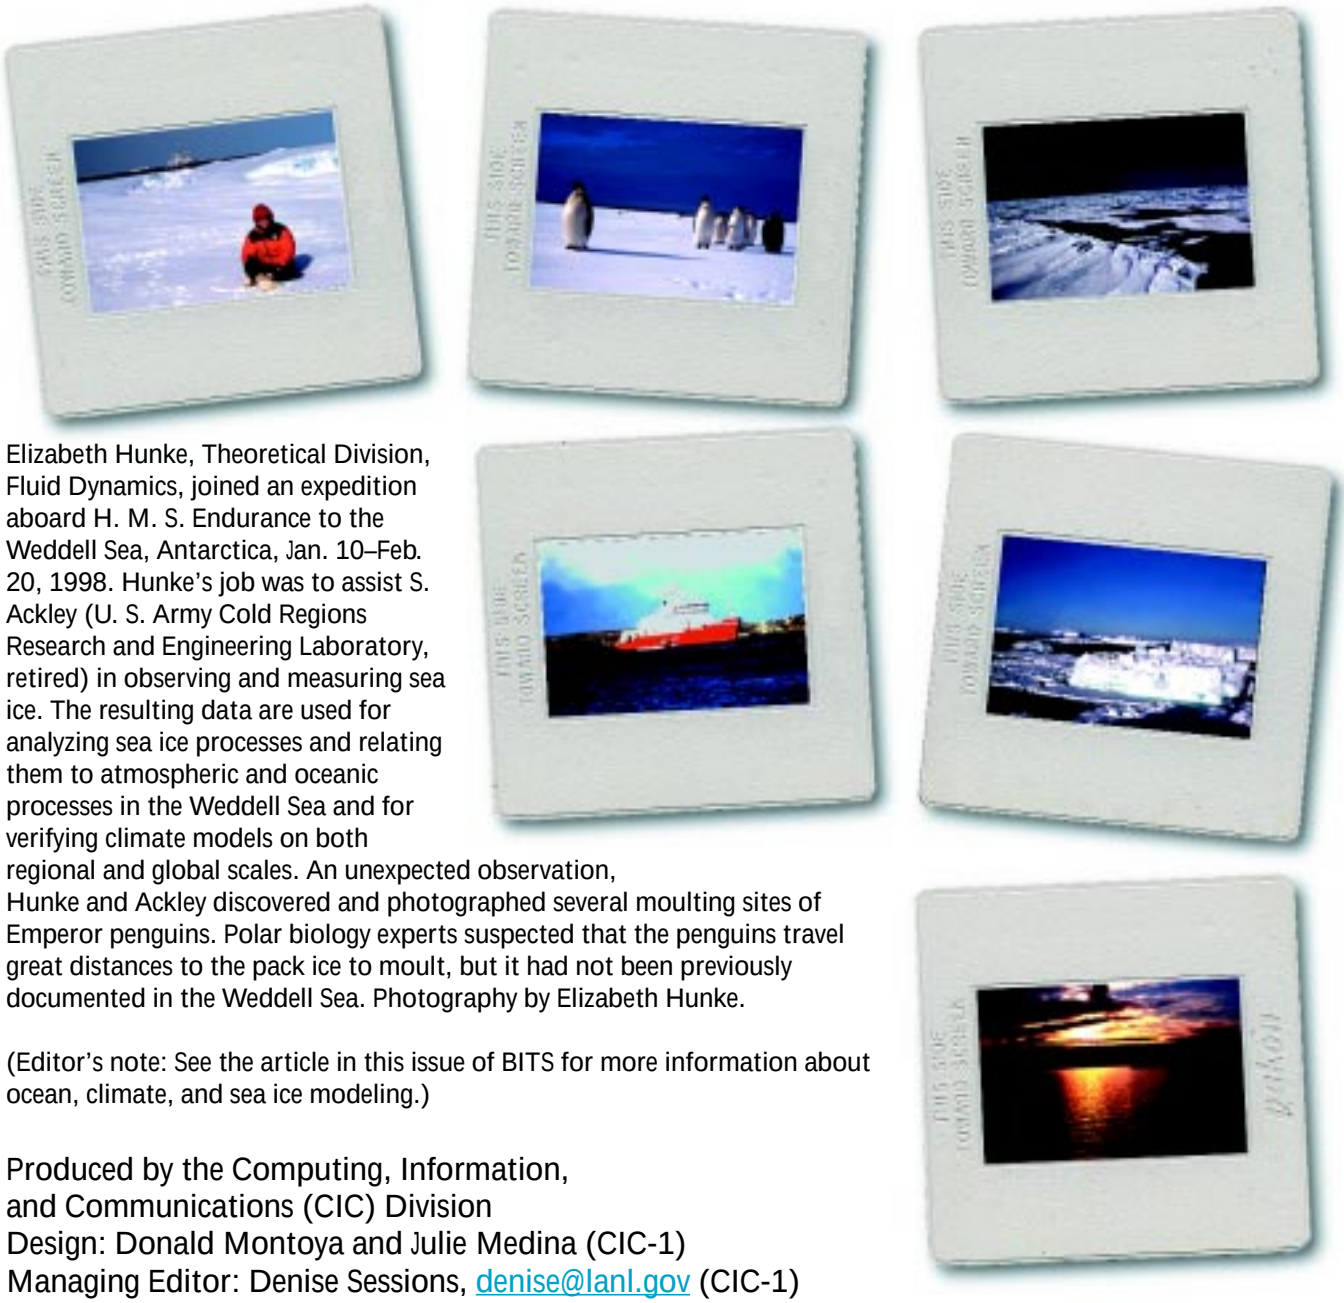

Elizabeth Hunke, Theoretical Division, Fluid Dynamics, joined an expedition aboard H. M. S. Endurance to the Weddell Sea, Antarctica, Jan. 10–Feb. 20, 1998. Hunke's job was to assist S. Ackley (U. S. Army Cold Regions Research and Engineering Laboratory, retired) in observing and measuring sea ice. The resulting data are used for analyzing sea ice processes and relating them to atmospheric and oceanic processes in the Weddell Sea and for verifying climate models on both regional and global scales. An unexpected observation, Hunke and Ackley discovered and photographed several moulting sites of Emperor penguins. Polar biology experts suspected that the penguins travel great distances to the pack ice to moult, but it had not been previously documented in the Weddell Sea. Photography by Elizabeth Hunke.

(Editor's note: See the article in this issue of BITS for more information about ocean, climate, and sea ice modeling.)

Produced by the Computing, Information, and Communications (CIC) Division
Design: Donald Montoya and Julie Medina (CIC-1)
Managing Editor: Denise Sessions, denise@lanl.gov (CIC-1)
Illustration: Dave Delano, Art/Illustration
Printing: Imaging Services Group (CIC-9)
HTML: David Van Etten (CIC-1)

Attention Subscribers: There are two ways you can update your subscription data. E-mail bitsupdate@lanl.gov, or fill out the form inside this issue. Internal LANL subscribers please note that until we have a mechanism in place to verify our subscription database with the LANL Employee Information System we will rely on you to update your mailstop. Thanks for keeping us updated on subscription data.

Table of Contents

Achieving Revolution through Evolution:

Breaking Giant Leaps into Small Steps

by John H. Hall, Technical Staff Member, X-10: Code Group B

The time scales for computing projects at Los Alamos are unprecedented in industry. Here, coding teams are often formed and maintained for decades. For a new project leader on a relatively new project in Laboratory time (three years old), deciding the best path forward can be rather daunting. Coding for the entire lifecycle of the product takes on new meaning when you are probably going to be the one maintaining that product. How do you balance demonstrating continuous progress against promoting design sufficient to prevent the eventual collapse of the entire system? In this month's article I am going to present some observations about how the Tecolote Team has approached these issues.

The Nightly Build

A doctor began having trouble with his car. He promptly took it to the local garage where the mechanic quickly diagnosed and fixed the trouble. As he was preparing the bill, the mechanic noticed that his customer was a doctor and commented, "You know, doc, you and I both do pretty much the same job. We diagnose a problem, remove the bad parts, and replace them with new ones." The doctor smiled and said, "yes, but I don't have the luxury of turning the motor off until I am done."

Although the Tecolote project is only three years old, it has been in use by various customers for some time now. An obvious issue in such a case is how long can we go without a working code during a major upgrade? Our conclusion was we could *never* afford to be caught without a working code. Evidence that we are not the only code team to arrive at that conclusion came in the number of tools that existed with options to help us achieve our goal of always having a working code.

First, the cvs repository software used in revision tracking had developed the idea of branches so that unchecked or malfunctioning modifications could exist on a branch and still be registered (taking advantage of the automatic backups of the repository). A branch belongs to the programmer making the changes who is in complete control when it is reintroduced to society. Tools exist for grafting the branch back into the main trunk when the software reaches a sufficient quality level. This is roughly the equivalent for software developers of shutting the engine off for a mechanic.

Second, the industry had coined a term for a series of practices that the Tecolote team has followed religiously for almost its entire project life. These practices are collectively known as the "nightly build." As soon as we had a working application developed within our framework, we automated the checking out each night of the "HEAD" version of our code from the cvs repository, which was then compiled and tested against a series of "regression" tests specifically designed to prevent the reintroduction of problems we had previously encountered. The formality of operations we developed to survive the continuous adoption of new methods while not breaking the chain of successful nightly builds has allowed us to code more like the doctor, who has to deliver while keeping the engine going.

At Microsoft, the software engineer that checks in code that breaks the nightly build becomes responsible for running it each night until someone else makes a mistake. Since we have automated the nightly build, this is not much of threat since it would probably only mean that you have to add a single line to your crontab file to get the entire process going each night. So we have resorted to the time-honored tradition of ridicule for the person who gets caught being sloppy. Peer pressure can be a particularly effective motivator.

The nightly build has become such a dominant piece of the Tecolote culture, that few people even use the branch constructs of cvs anymore (which can sometimes be untrustworthy due to complex interactions between tens of programmers). We develop and continuously test as we are going along, and at any point in which all the regression tests are successfully passing, we are free to check in.

The Logistics of the Campaign

In the movie Patton, George C. Scott spends considerable portions of the dialog explaining how it is the chief job of the commanding general to ensure that logistical supplies necessary for the successful completion of the campaign are in place at precise times throughout the campaign. These of course include food, personnel, armored divisions, and most importantly, fuel.

Doctors follow a similar methodology in preparation for major surgery. They often practice complex procedures as an assembled team to ensure that at the appropriate time, the correct instruments are present so that continuous progress can be made.

How do these practices reveal themselves in software development? First, if you are going to prevent any interruptions in the nightly build process you need to isolate the interactions of these new, untested parts with the rest of the complex system. This is similar to closing off blood flow to an appendix before removing it by clamping down on the blood vessels connected to it. Fortunately, the software industry has made considerable progress in this area with the introduction of object-oriented programming. I will present below an example of how coding practices can be changed to anticipate future modification and minimize the impact throughout the code. This design pattern is one of the most significant developments of object-oriented programming.

Second, interruptions in the workflow must be identified and rooted out during the planning phase. Since anyone who has programmed on a real project knows that it is impossible to envision all the complex interactions in advance, a slightly different tact must be taken to prevent workflow interruptions. The job of the software architect changes from predicting all possible interactions in advance, to preventing the premature closing of an alternate path. In other words, the designer should not intend to identify and follow a single path (the ultimate form of arrogance), but should plan instead to keep as many options open as possible throughout the campaign. This “designing for flexibility” is a hallmark of the Tecolote project and influences every decision we make.

Finally, we must keep our objectives simple so that they can be quickly explained to others. In the movies, you can always recognize the well-trained SWAT team or SEAL team because they have paint on their faces and are wearing dark clothing (the jackets labeled SWAT sometimes help) and because their leader is constantly making a series of hand motions which are describing complete preplanned courses of action. As a team reaches maturity, the past history together allows a more complex dialog rooted in shared experiences. Constant vigilance must be maintained to allow newcomers on the project to become members of the team. Exclusionary practices such as references to obscure events prevent the newcomer from coming up to speed as quickly as possible. Perhaps we need a programmer “Boot Camp” to help them develop a common dialog and learn some of the common hand-signals.

Getting it Right from the Beginning

In procedural languages such as C or FORTRAN, the decision flow is maintained through state variables or flags. Every decision that might vary depending on the state of a given flag would require an if test wrapper either at compile time using a preprocessor #if or at run-time using an actual conditional expression. A consequence of this hardwiring of decisions throughout the source code base is that when a change is needed, each of these conditional expressions must be reexamined and modified to ensure the correct logic still obtains. This rapidly becomes unwieldy and the source code becomes impossible to maintain.

One of the major benefits of object-oriented programming derives from eliminating this propagation of conditional testing by replacing the myriad if tests with a single test that determines which object is to be constructed. That object is built to carry out only one possible path through the conditional forest and different objects are constructed for each of the other possible paths (if needed).

A simple example involves the parallel reading of a simple text-only input file. Opening the same file thousands of times (once for each processor) is a massive waste of resources. A much better technique is to have processor 0 open and read the file and then broadcast the information to the other processors. The following sample code shows an inheritance hierarchy to accomplish this task.

```
class Redirect {
public:
  virtual int get_char(void)= 0; // Pure virtual function
};
class RedirectFile : public Redirect {
public:
int get_char(void) {
  // Reads file into buffer and broadcasts result.
  // Returns a single character from the buffer at a time.
}
};
class RedirectInput : public Redirect {
public:
int get_char(void) {
  // Gets buffer information from processor 0. Also returns
  // a single character at a time.
}
};
int main (int argc, char **argv) {
// Initialize parallel communications
.
.
.
Redirect *input = 0L; // Pointer to base class to hold new object
// Are we processor 0?
if( Pooma::myNode() == 0) {  // Yep, we are processor 0.
  input = new RedirectFile();
}
else {  // Not processor 0
  input = new RedirectInput();
}
// Use the Redirect object to accomplish your task
while ( input->get_char() != EOF) do_something();
return 0;
}
```

Using this technique, we have isolated the decision as to whether we belong to processor 0 to one place. All subsequent actions will occur correctly independent of the processor on which we actually reside. If at a later time, we need to be fed data through an in-memory communication, all we have
to do is derive a new class to receive the communication on processor 0, and then rebroadcast that message to the other processors as before. Then by editing the single if test in main we can install the new object and we are done. By writing a slightly more complicated if test we can choose which method to use dependent on an input parameter. In the object-oriented programming paradigm, almost all if tests should be replaced by this object construction pattern. This makes code more maintainable and extensible since one need only create a new class and change one place in the code to alter its behavior throughout. By replacing the if test with an abstract factory and a parser, new functionality could even
be added through a dynamic link library years after the main code was compiled and linked. This flexibility is crucial to the Tecolote Framework.

Summary

To be able to achieve revolution through evolution requires more and different planning than simply trying to accomplish a single focussed objective. The revolutionary designer must concentrate on preventing obstacles to future, yet unknown objectives from preventing necessary growth and redirection of the project. This concentration on keeping options open rather than finding a single path through the coding maze is difficult and requires new software designs. While these designs are sometimes more complex than the patterns they replace, maintenance is greatly reduced and significant changes to the infrastructure can often be accomplished while isolating these changes to a single place in the code.

Acknowledgment

I would like to acknowledge the contributions of Mark Zander, the chief architect of the Tecolote Project who taught me the above techniques during several years of protracted debate. These designs have withstood the test of time and have resulted in a framework that has already been extended into several different application areas without having to be completely redesigned. This flexibility is a tribute to Mark's vision.

Editor's note: BITS is featuring a series of articles on Tecolote. These articles have been published in previous issues, which can be accessed online at http://www.lanl.gov/orgs/cic/cic6/bits/archive.html .

- *Tecolote: An Object-Oriented Framework for Hydrodynamics Physics*, August 1999.
- *Using Tecolote Components to Extend Object-Oriented Programming*, September 1999.
- *Frameworks are Models, Too!*, October 1999.
- *Dimensionless Coding Techniques in Tecolote: Using a Single Source-Code Base for Multidimension Programs*, November 1999.

Ocean and Climate Modeling

by Rick Smith, Technical Staff Member, Theoretical Division, T-3 Fluid Dynamics

The Earth's climate is controlled by the complex interaction of many physical systems including the atmosphere, the ocean, sea ice, the land surface, and the biosphere. To be able to predict future climate change, it is necessary to understand both the natural variability of the climate system and the extent to which human activities affect it.

The ocean is of key importance in understanding climate, because changes in ocean circulation patterns are believed to be of primary importance in controlling climate variability on timescales of decades to centuries.

The Ocean's Mesoscale "Eddies"

Global ocean modeling is considered a Department of Energy High Performance Computing and Communication Grand Challenge problem both because of the societal importance of understanding climate variability and because it poses a severe computational problem. Realistic global ocean simulations are very difficult because the ocean contains both very small spatial scales and very long time scales. Most of the kinetic energy in the ocean is contained in the so-called "mesoscale" eddies, whose sizes range from 10 to 300 km. These eddies constitute the "weather" of the ocean. They are the oceanic equivalent of high and low pressure systems in the atmosphere, where the spatial scales are much larger: weather fronts typically extend over distances of 1000 to 3000 km.

On the other hand, the time scales in the ocean are much longer than in the atmosphere. Temperature anomalies in the atmosphere persist for at most a few months (unless they are associated with longer-term anomalies in the ocean surface temperature, as occurs in an El Niño event). The ocean, due to its inertia and large heat capacity, has a much longer memory: water mass properties in the deep ocean can reflect conditions that existed at the surface hundreds of years in the past. Residence times of deep water masses are typically several hundred years, and the ventilation time scale for the deep Pacific Ocean is more than a thousand years. Because of this, the integration time required to spin up an ocean model to a quasi-equilibrium state is several thousand years.

Feasibility of Eddy-Resolving Climate Simulations

The computational resources required to integrate a global ocean model with spatial resolution of order 10 km for 1000 years or more will not be available for the foreseeable future. Current ocean models used in climate simulations have spatial resolutions ranging between 100 and 400 km. High-resolution basin and global-scale models with order 10 km resolution can only be integrated on today's supercomputers for a few simulated decades. Thus another factor of 100 increase in computing power is needed before eddy-resolving climate simulations become feasible, and it will likely be a decade or more before such resources become available.

What To Do in the Meantime?

The Climate, Ocean, and Sea Ice (COSIM) project at Los Alamos is attacking the climate problem from two different directions. First, in collaboration with the National Center for Atmospheric Research (NCAR) we are developing coupled climate models using low- to moderate-resolution ocean components. Two coupled-climate models at NCAR—the Climate System Model (CSM) and the Parallel Climate Model (PCM)—have adopted the Los Alamos Parallel Ocean Program (POP) and the Los Alamos Sea Ice model (CICE). More information on climate, ocean, and sea ice modeling at Los Alamos is available on our Web server: http://www.acl.lanl.gov/climte.

The other approach we have taken is to make use of the supercomputing resources at Los Alamos to conduct very high-resolution eddy-resolving ocean simulations, albeit of shorter duration, using the POP model. There is much that can be learned from these shorter high-resolution simulations. Ten to 30 years is sufficient time for the model to reach a quasi-equilibrium state where the velocity field has adjusted to the initial density field, which is determined from observed distributions of temperature and salinity. These short simulations are therefore appropriate for studying the dynamics of the ocean circulation given its observed density structure on time scales of less than a decade, but not for studying long-term evolution of deep water masses or climate variability on time scales of decades and longer.

Nevertheless, the high-resolution simulations are very important for climate research, because the model output provides quasi-realistic fields of turbulent statistics that can be used to guide the development of subgrid-scale parameterizations for use in coarse-resolution climate simulations. Furthermore, the model provides comprehensive three-dimensional datasets that can aid in the interpretation of the extensive observations taken over the last decade, such as high-quality satellite altimetry measurements and the variety of *in situ* measurements collected as part of the World Ocean Circulation Experiment (WOCE).

The North Atlantic Ocean at 1/10° Resolution

The first major simulation performed with the POP model was a global ocean simulation driven by observed surface winds for the decade 1985 to 1995 (Maltrud et al., 1998). This model had a horizontal resolution of 0.28°, which was sufficient to allow the development of an eddy field, but not sufficient to fully resolve the spectrum of mesoscale turbulence. These and similar "eddy-permitting" simulations conducted by other researchers were able to reproduce many aspects of the wind-driven circulation. However, they were unable to reproduce some basic features of the mean circulation, such as the separation points of the Gulf Stream and other western boundary currents and the observed paths of major current systems like the North Atlantic Current (which flows along the Grand Banks east of Newfoundland).

Detailed analysis of the global simulation compared to satellite observations (Fu and Smith, 1996) clearly demonstrated the need for even higher spatial resolution, and theoretical arguments suggested a horizontal resolution of 1/10° to 1/12° would be needed. At that time a global simulation was not feasible at this resolution, so we opted to conduct a limited-domain simulation of the North Atlantic Ocean at 1/10° (Smith et al. 1999). This model is also driven by observed winds covering the period 1985 through 1998. The model domain extends from 20S in the South Atlantic to 72N, and includes the Gulf of Mexico and the western half of the Mediterranean Sea.

The figure on the cover of this issue shows a snapshot of the sea surface temperature from the North Atlantic simulation, with red colors indicating warm and blue colors indicating cold water. The path of the Gulf Stream, which carries warm water from the tropics to high latitudes, can be clearly seen. The current follows the coastline of the southeastern US until it separates from the coast at Cape Hatteras; from there it begins to meander and pinches off warm and cold core eddies. In the 0.28° global simulation (and in other simulations of comparable resolution) the Gulf Stream does not separate at Cape Hatteras as observed, but continues much farther north before separating; this has been a long-standing problem with ocean circulation models.

Eddy Variability

A remarkable feature of this simulation is the emergence of a ubiquitous mesoscale eddy field which is substantially stronger than seen in previous simulations. The eddy kinetic energy constitutes 70% of the total basin-averaged kinetic energy. The model results agree well with observations of various different measures of the eddy variability, including the magnitude and geographical distribution of near-surface eddy kinetic energy and sea-surface height variability, the wave number-frequency spectrum of surface height anomalies in the Gulf Stream, and measurements of the eddy kinetic energy as a function of depth. Thus the model appears to be close to capturing the bulk of the energy in the mesoscale eddy spectrum.

Figure 1 shows the root-mean-square surface height variability from the model averaged over a 4-year period (a), and from a recent high-quality blend of altimeter data from the TOPEX/POSEIDON and ERS satellites (b). This level of agreement between model and observations is unprecedented, and represents a milestone for both numerical ocean modeling and satellite altimetry.

Because of the geostrophic balance between the Coriolis force and pressure gradients, changes in the sea-surface height reflect the near-surface ocean currents. Mesoscale turbulence is generated by instabilities of the mean flow, and hence the eddy field is most intense in regions of strong western boundary currents. This can be seen in the figure: the height variability is most intense in the region of the Gulf Stream extension (around 30N to 45N latitude and -75E to -50E longitude) and in the vicinity of the North Atlantic Current (40N to 50N and -50E to -35E). Some regions of high variability that appear in the observations but not in the model (such as off the west coast of South America near the equator and off the North American coast southwest of Nova Scotia) are residual errors associated with the removal of tides from the altimetry measurements. Time series of two-dimensional fields of the model surface height are now being used by scientists in the US and France to help interpret the existing satellite altimetry measurements and develop the next generation of satellite altimeter experiments scheduled to be flown in the next few years.

Mean Circulation

While the agreement between the eddy variability in the model and observations is impressive, what is most remarkable about this simulation is that the time-mean circulation exhibits several significant improvements relative to previous simulations: 1) the Gulf Stream separates at Cape Hatteras, and its peak velocities, transports, spatial scales, and cross-stream structure are in good agreement with current-meter data; 2) south of the Grand Banks the Gulf Stream splits into the northeast-flowing North Atlantic Current and a southward flow that feeds the Azores Current; 3) the time-mean path of the North Atlantic

Current is in good agreement with observations from float data, including the detailed positions of troughs and meanders; and 4) this is the first realistic simulation which correctly simulates the Azores Current, which flows eastward at about 35N in the central and eastern basin. Its position, total transport, and eddy variability are consistent with observational estimates. The surface-height variability associated with this current can be seen in Fig. 1 as a tongue of high variability between 33N to 35N and -40E to -20E that appears in both the model and observations.

This simulation is by no means perfect. There are notable discrepancies with observations in some areas, for example, the Gulf Stream separates at Cape Hatteras, but its eastward path after separation is about 1.5° too far south. Nevertheless, the overall improvement in the time-mean flow relative to previous simulations indicates that we have crossed a threshold in resolution and entered a new regime of the flow that is much closer to the real circulation of the North Atlantic. This strongly suggests that the turbulent eddy field plays a crucial role in determining the character of the mean flow. We are in the process of diagnosing statistics for eddy-mean flow interactions to understand more precisely how the eddies affect the mean flow and to guide the development of turbulence closure schemes for use in coarse resolution climate models.

What's Next

Spurred by the success of the North Atlantic simulation, and with encouragement of ocean modelers and observational oceanographers throughout the community, we are currently making plans to run a 1/10° global simulation. If the computing resources become available, we hope to begin a 10- to 30-year simulation

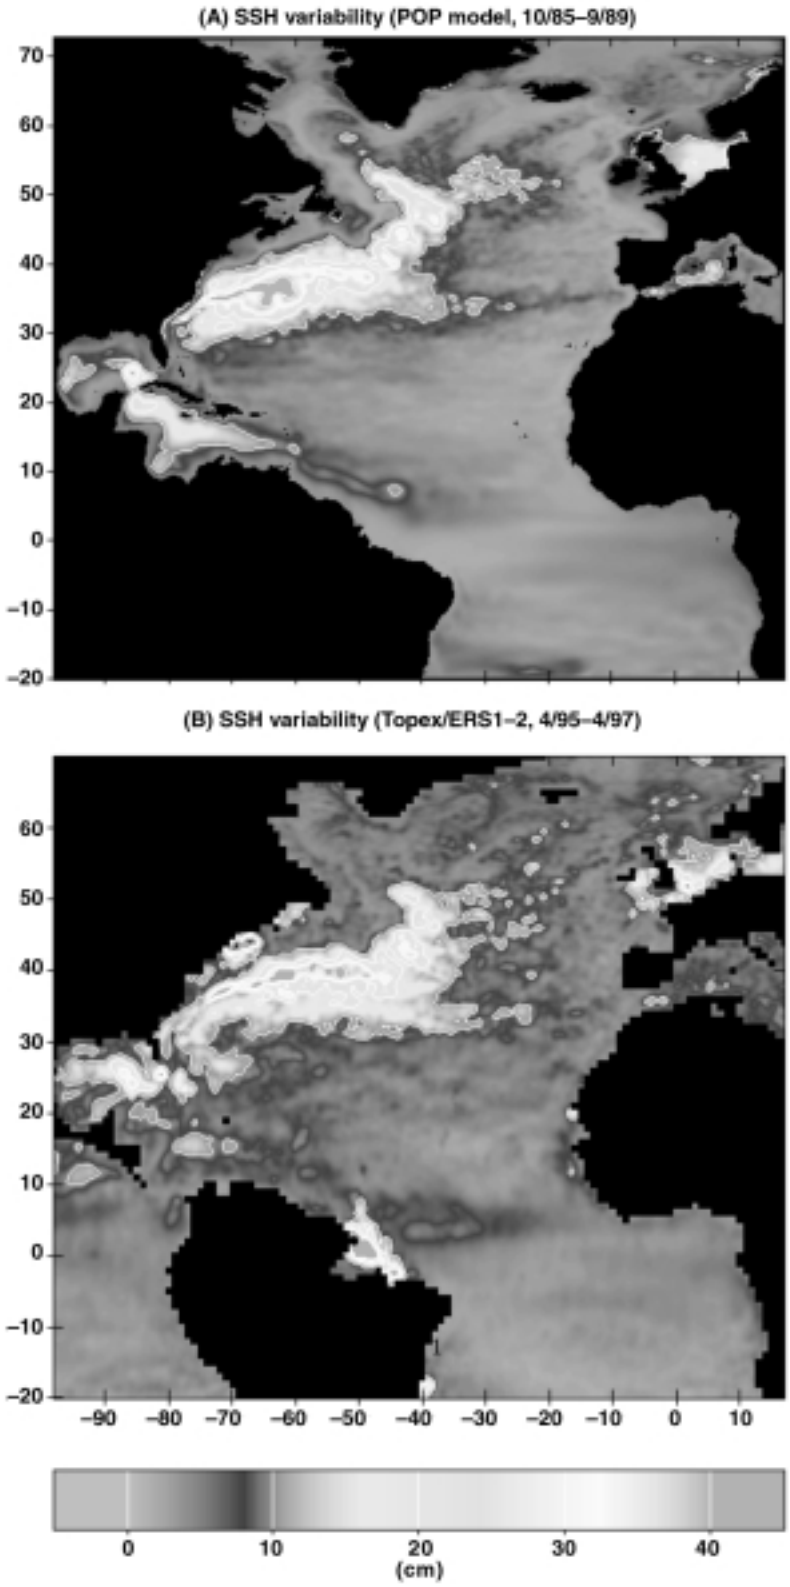

Fig. 1. Root-mean-square sea surface height variability in the North Atlantic Ocean: (a) Los Alamos POP model simulation (Smith et al., 1999); (b) altimeter observations in a blended analysis of data from the TOPEX/POSEIDON and the ERS 1 & 2 satellites (Le Traon and Ogor, 1998; Le Traon et al., 1998).

this spring. This will be by far the most ambitious ocean simulation ever undertaken, but the scientific payoff will be large if the global model shows the kind of improvements seen in the North Atlantic model.

For more information on POP, CICE, and climate modeling at Los Alamos see http://www.acl.lanl.gov/climte .

Editor's note: The graphic on this issue's front cover was produced using the POPTEX visualization program developed by Alan McPherson in the Advanced Computing Laboratory. For more information see http://www.acl.lanl.gov/viz.

References

L.-L. Fu, and R. D. Smith, "Global Ocean Circulation from Satellite Altimetry and High-Resolution Computer Simulation," *Bull. Amer. Met. Soc.* **77**, 2625–2636 (1996).

P. Y. Le Traon, and F. Ogor, "ERS1/2 Orbit Error Improvement Using TOPEX/POSEIDON: The 2-cm Challenge," *J. Geophys. Res.*, **103**, 8045–8057 (1998).

P. Y. Le Traon, F. Nadal, and N. Ducet, "An Improved Mapping Method of Multi-Satellite Altimeter Data." *J. Atmos. Ocean. Tech.*, **15**, 522–534 (1998).

M. E. Maltrud, R. D. Smith, A. J. Semtner, and R. C. Malone, "Global Eddy-Resolving Ocean Simulations Driven by 1985–1995 Atmospheric Winds," *J. Geophys. Res.*, **102**, 25203–25226 (1998).

R. D. Smith, M. E. Maltrud, F. O. Bryan, and M. W. Hecht, "Numerical Simulation of the North Atlantic Ocean at 1/10°." *J. Phys. Oceanogr.*, in press (1999).

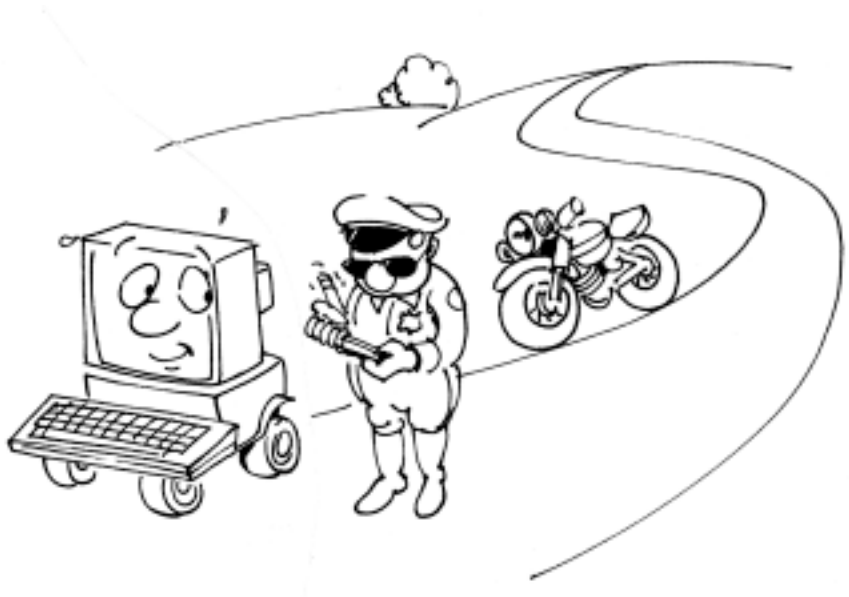

New Electronic Resources Available at the Research Library

New Version of Science Server® at LANL

by Miriam Blake, Databases Team Member, CIC-14, Research Library

As of November 1, 1999, a new version of Science Server® at LANL http://sciserver.lanl.gov is available.

What is Science Server® at LANL? Science Server® at LANL is an electronic product created to promote scientific communication and collaboration. Science Server® at LANL manages content and delivers electronic journals directly to your desktop.

The Research Library has created Science Server® at LANL to provide a single solution for searching, browsing, and delivery of search results through a standard Web browser interface.

Science Server® at LANL consists of three pieces: the electronic content, the content management software, and the value added by integrating content in a single package.

The electronic content provided by Science Server® at LANL currently includes 1100+ scientific, medical, and technical journals from publisher Elsevier Science.

Current Version

Science Server® at LANL now runs on a locally modified version of the 3.0 software provided by Science Server LLC. New features in this version include:

- Quick title searching on the main page,
- Improved navigation, and
- Substantial changes to the architecture, including dynamic page generation and improvements to the indexing of new articles.

Future Plans

Plans for future additions to Science Server® at LANL include:

- IEEE—approximately 120 journals, more than 600 conference proceedings titles, and over 875 technical standards—published since 1988
- Academic Press, 128 journals
- Kluwer, 365 journals
- Wiley, 74 journals
- IOP, 34 journals
- ACS, 30 journals
- Springer-Verlag, approximately 450 journals

Science Server® at LANL is accessible at http://sciserver.lanl.gov/ or via the link in the middle of the Electronic Journals page.

New Engineering Index® at LANL

by Kathy Varjabedian, Databases Team Member, CIC-14, Research Library

A new interface and search engine is now available for Engineering Index® at LANL (http://engindex.lanl.gov/lanl/), our version of the world's most comprehensive interdisciplinary engineering database. The new system is similar to SciSearch® at LANL.

Features include:

- Alerts—weekly e-mail delivery of new information for searches of your choice
- Easy limiting by date, language, document type, and electronic format
- Phrase searching
- Sorting of results by relevancy ranking, date, author, or title
- Ability to search individual author names with initials
- Links to both print and electronic research library holdings
- Ability to Mark All for printing, e-mail, or downloading

Engineering Index contains over four million abstracts. Each year, over 220,000 new abstracts are added from almost 4,500 international journals, conference proceedings, books, technical reports, and dissertations. Coverage spans over 175 disciplines and major specialties within engineering. Engineering Index corresponds to the print publication by the same name. Coverage begins in 1969 and is updated weekly.

Ghost Helps CIC-2 Back Up Data

by Randy Cardon, CIC-2 Managed Desktop Project Leader, and Nikki Gaedecke, Information Architecture (IA) Desktop Team Member, IA Project

Third Article in a Three-Part Series

Many people who lose data can revert to backup files. For individuals who fail to back up important documents, they risk being forced to recreate them should an "accident" occur. Imagine coming into work one morning to find everything on your computer is gone: no applications, no files, no Word documents, and no projects that you've been working on for the last couple of months. Gone, and you don't have your data backed up!

While Halloween has come and gone, something scarier haunts computer support personnel: the Year 2000 bug.

In previous BITS publications (August and November 1999), we focused on several reasons why backing up data is important: it is a DOE requirement, it makes economical common sense, and it provides insurance against unforeseen computer and software failures. To back up data, end users and system administrators can choose between tools such as floppy disks, removable media, area network servers, or Adstar Distributed Storage Manager (ADSM). Choosing the most appropriate backup utility, depends on the user's amount of data, frequency of backups, budget, involvement (manual vs. automatic backups), type of data (classified vs. unclassified), and the computing environment itself. The following rule of thumb depends on your bottom line: how often you should back up your data depends on how much you can afford to lose.

Using Disk Partition Imaging

In the advent of Year 2000, the Desktop Support Group (CIC-2) is taking a different approach to backups and more specifically disaster recovery. If for some reason you lose everything on your computer because of Year 2000, a disk crash, or even a fire, and you have your data safely backed up, CIC-2 technicians can quickly help their customers recover using Norton Ghost™, a disk partition imaging product. Unlike removable media and ADSM, Norton Ghost enables computer technicians to completely restore a computer in 30–45 minutes instead of 3–4 hours.

By using Norton Ghost, Randy Cardon, Project Leader for the CIC-2 Managed Desktop Project, and his team can provide better computer support by standardizing the file structure, which is a stepping stone for remote administration. Before Norton Ghost and the managed desktop philosophy, CIC-2 computer technicians would have to spend over three hours to rebuild one crashed computer. To make matters worse, technicians were and are still finding that each serviced computer would have different setups and file structures. Under these circumstances, technicians would be confronted with challenging software upgrades and would have to "chase-down" data files that could be missed when backing up data. With too many machines, too few computer technicians, and no standards for file structures, CIC-2 formed the Managed Desktop Project to confront these problems using Norton Ghost as one tool.

Beginning with Microsoft® Windows NT® and Macintosh®, Randy Cardon and his team are working on disk images of standardized configurations. Analogous to the Mac, they install Windows NT, all the service packs, hot patches, a set of applications, and utilities. Then, using Norton Ghost they create an image (or duplicate) of this whole environment. The image can be stored on a network drive, CD, or Jazz drive, for example. Then, using a bootable floppy or CD, they move the image to the stripped-down computer. After booting the computer in Windows NT, they change the security identifier (SID), add the IP address, configure the network and video card, and add user information (i.e., Eudora® settings).

Standardizing File Structures

In addition to backup and disaster recovery options, the Managed Desktop Project team hopes to standardize computer file structures to pave the way for remote administration. For example, software such as Systems Management Server (SMS) can perform software upgrades overnight, but unfortunately, nonstandardized file structures can prevent the remote upgrade to perform smoothly.

For More Information

If you would like more information about the Managed Desktop Project, e-mail Randy Cardon at rec@lanl.gov . Norton Ghost is available on the Electronic Software Distribution (ESD) Web site at http://esd.lanl.gov . For general questions about back up methods, contact your system administrator or the IA Desktop team at ia-desktop@lanl.gov .

"IT'S A LAPTOP MAINFRAME."

Fast Reduction of Analytical Data Using WINDOWS® 95/98/NT

A Software Application for ICP Mass Spectrometry

by John R. Quagliano and Deborah J. Figg, Technical Staff Members, NMT-1 Analytical Chemistry, Nuclear Materials and Technology (NMT) Division

Modern analytical instrumentation is capable of producing large data sets in a short amount of time due in part to multichannel detection and computer automation. However, it is still the analyst's responsibility to interpret and organize the raw data into a useful and meaningful format and also to assure that numerous Quality Control (QC) criteria are satisfied. It has been our experience in NMT Division that instrument vendor and third-party software programs are inadequate for these data reduction needs. Consequently, our chemists spend an inordinate amount of time performing spreadsheet manipulations to complete the data reduction and there is much less time spent in the lab developing improved measurement techniques. As a result, the analytical services team produces a slower turnaround of results to customers.

To solve this problem, we have developed a stand-alone executable software program that reformats raw data, performs QC checks, notifies the analyst of errors, prepares data for a Laboratory Information Management System (LIMS), and generates customer deliverable reports. The program FRAD™-MSanalyzer is mouse/menu driven with the familiar Windows® Graphical User Interfaces (GUIs) and is suitable for technicians as well as scientists or engineers. As an example, this short article will demonstrate how a 6-isotope Inductively Coupled Plasma Mass Spectrometry (ICP-MS) raw data file that previously required 2 hours to process should now take less than half an hour to process without any loss of thoroughness or quality of the analysis.

Determining Impurities in Inorganic Materials

At the Chemistry and Metallurgy Research facility and at Technical Area- (TA-) 55, NMT Division performs a variety of analytical activities in the chemical sciences to determine the levels of impurities in inorganic materials. The impurity of these materials in turn has a great bearing on the reliability of the materials for various intended uses. Some of the more common analytical techniques include Ion and Gas Chromatography, Radiochemistry, x-ray Fluorescence Spectrometry, Atomic Emission Spectroscopy, and ICP-MS. Many of these techniques are complementary, meaning that if one technique is incapable of identifying and quantifying a specific elemental impurity at the required trace levels, one of the other techniques will "come to the rescue" and pick up that analyte. However, ICP-MS is usually relied on for the most comprehensive and difficult analyses because of its superior detection limits, ability to measure isotopes, and wide coverage of the Periodic Table of Elements. Although data from an ICP-MS analysis are used as an illustration in this article, plans are underway to adapt the customized software to handle data from a variety of analytical instruments.

Development Tools and Plan

The source code is comprised of two modules. The bulk of the functions have been coded into a core engine that provides a common platform that is nearly independent of both the instrument type and the particular format of raw data file as generated by the vendor data acquisition software. The other module can be thought of as a cover for the engine that is dependent upon both the instrument type and the format of raw data file. This separation of functionality enables us to enjoy significant code reuse as only the covering module needs to be adapted on a case-by-case basis depending whether we are, for example, performing chromatography versus mass spectrometry. The covering module is essentially just a very large switch-case statement that performs ASCII text full-line recognition and sorting based on targeted character sequences in the data file.

We have used the Borland C++Builder™ Integrated Development Environment (IDE) version 3 from Inprise, Inc. to create the software described in this article. The IDE is a full fledged visual C++ compiler that produces code that only runs on the Microsoft® Windows® operating systems. To fully utilize the visual toolboxes, vendor implemented extensions to ANSI C++ are required (which is also true of the Microsoft® Visual C++ product). As a convenience to the end user, the executable has been statically linked with one redistributable Borland library and

therefore no installation or modification of the registry is required. As such, the executable, library file, and required input text files (requiring under 10 MB) can be run straight from portable media. A multitab sheet dialog box help is available during program execution as well as an online user's manual in Microsoft® Word® format. The software has passed a rigorous Quality Assurance (QA) validation and verification process per NMT Division requirements. A more complete listing of the primary tasks performed are given below, and the remainder of the article will briefly describe GUIs related to three of these tasks. We conclude with some statements about leveraging the current software to perform more advanced multivariate statistical analyses in the future.

Primary Software Tasks

These are the primary tasks performed by the software application

- reformats raw data file from instrument into an organized spreadsheet template,
- applies analyst-specified experimental factors such as degree of dilution,
- rounds raw numbers per defined precision rules,
- calculates QC values for instrument bias and drift trending,
- generates customer report interactively with Analyst, and
- saves all QC exception reports to disk.

Experimental Results

The unedited raw data file as acquired from the ICP-MS via the instrument vendor's data acquisition software is often formatted awkwardly with patterns of delimiters and control symbols that would be difficult to analyze as is. FRAD-MSanalyzer uses the covering module to reformat the contents, without altering their representation, into a readily viewable spreadsheet as shown in Fig. 1. This figure also shows part of a full-sentence exception report generated by the program during a menu-choice-invoked call to the QC checking functions.

Quality Control

No measurement can give a 100% completely accurate and precise number due to both random and systematic errors. In analytical chemistry, we strive to ensure that these two types of errors steadily fall within a predetermined acceptable tolerance range and this is what is meant by QC. Figure 2 illustrates a typical QC spreadsheet of values for Calibration Blank samples (ICB and CCB), Laboratory Control Standards (LCS) and Serial Dilution (Dilution) samples. The appropriate cells are populated with isotope names, sample names, and control values automatically by the program, which in addition to saving time also eliminates manual transcription errors. Failing values are marked with an asterisk for quick identification. Note how the LCS values, though passable as prescribed by the ±20% QC limits, are all biased slightly high versus the 50 ppb standard defined as 100(%). Over a period of weeks this result can be trended to see if a fundamental drift (systematic error) has developed in the mass spectrometer's response to these isotopes. This is of immense value to the chemist, and, previous to FRAD-MSanalyzer, there was no convenient way of control charting our ICP-MS results.

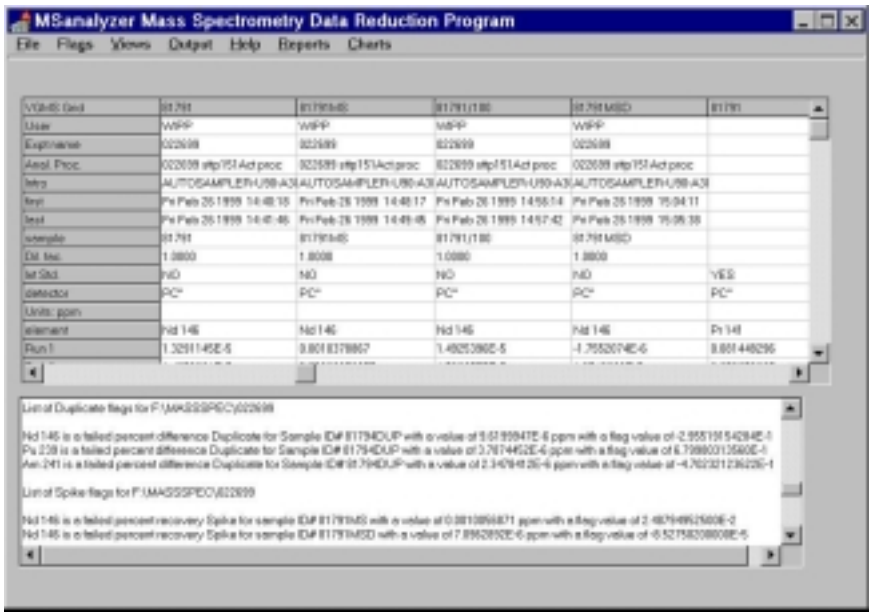

Fig. 1. Reformatted Data and Quality Control Exception Report

Data Reporting

The final deliverable of an analysis of a sample is a spreadsheet report to the sample submitter. Previously, the analyst spent a significant amount of time manually transcribing concentrations, running copy and paste operations, and visually inspecting or hand calculating variables in Microsoft®

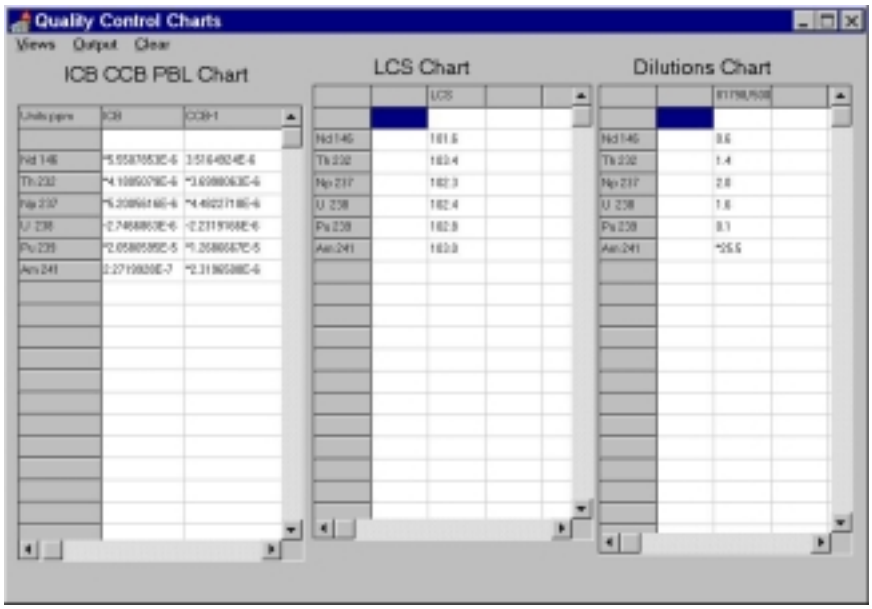

Fig. 2. Quality Control Chart values for Blanks, Standards, and Dilutions

Excel®. The task is now completely mouse driven; Borland library routines use appropriate calls to Windows® API functions that permit mouse-to-spreadsheet coordinates to be simultaneously synchronized. Figure 3 shows a partially processed report window as manipulated by the end-user. In this example the analyst has applied subject matter experience to choose between reporting values from either the more or less concentrated of two partnered samples. In the case of Am 241, the "G" flag in front of 0.0017 for sample 81798/500 signifies a serial dilution failure; the value of 0.013 is chosen instead. The tagged values on the left grid are transferred to the right grid with a simple menu selection. The result is a visually oriented and fast user selection of a subset of values that actually go into the final report are displayed on the left. Once committed, the report on the right is saved as a comma separated value (.csv) file readily viewable in Excel® or a variety of databases.

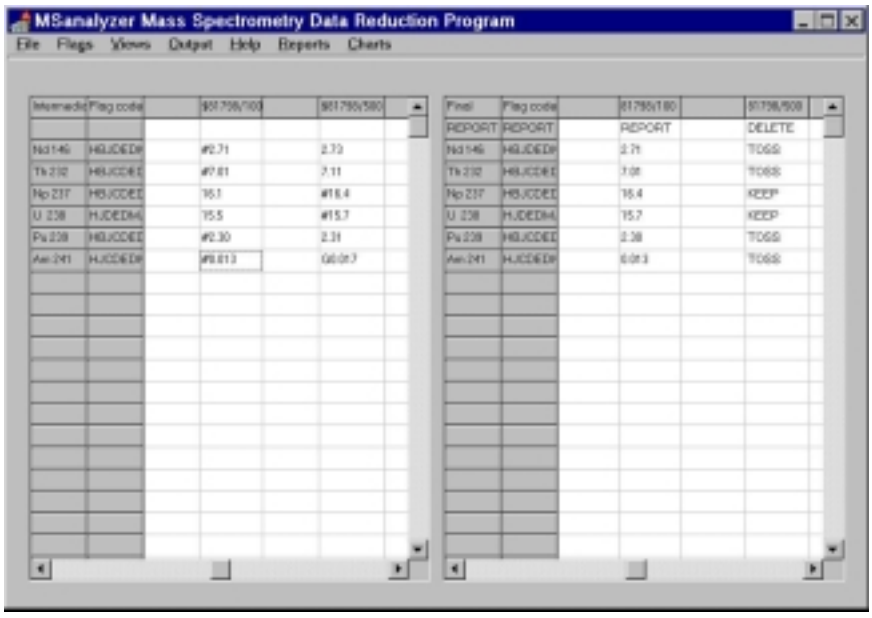

Fig. 3. Customer Deliverable Report Spreadsheets

The alert reader will note that we only addressed how to conveniently produce one of NMT-1's common report deliverables – the customer report. In other words, what about the upload to LIMS? Indeed, we have nearly completed QA testing of a completely distinct program that will create files easily for dependable upload to LIMS. This second program will pick up the output of FRAD-MSanalyzer and perform the necessary conversion with some user interaction, roughly analogous to the passing of a baton between runners in a relay race in a track and field event. Once completed, it is our intention to submit a report on this second program in a future BITS publication.

Future Activities

Our ultimate goal is to have this conveniently formatted set of data results serve as input for the commonly used and already available chemometric software packages. Some of the statistical techniques which have proven useful include Partial Least Squares Regression, Principal Component Analysis, and K-nearest neighbor (no need to include acronym unless you use the term again) clustering, just to name a few. In this way, the combined FRAD™/Chemometrics software can make the basic chemometric tools more attractive to analytical service personnel and assist R&D personnel in assessing the quality of their research data before treatment by advanced chemometric modeling tools. We hope to also report on these activities in more detail in a future BITS article.

New Tool for System Administrators of LANL NT MAD Accounts:

Domain Manager Application

by Sheila T. Molony, Technical Writer/Editor, CIC-1 Communication Arts & Services

The number of user accounts on the LANL Windows NT Master Accounts Domain (MAD) is growing into the thousands, and NT network administrators now have a new tool to create and manage user groups and accounts more easily.

The new Web-based application is the LANL Domain Manager. Please contact Cheryl Host (chost@lanl.gov) or Bob Stewart (res@lanl.gov) to gain access to this application.

Background Information on LANL MAD

The LANL MAD is the central domain for all Windows NT user accounts. Individual users can register for accounts in the LANL Domain at http://register.lanl.gov. Administrators can migrate users in existing domains to the LANL Domain by a number of methods. To avoid re-entering each account contact Cheryl Host at chost@lanl.gov for migration assistance. For an introductory article about LANL MAD see Cheryl Host's September 1999 BITS article—"News About Microsoft Networking at Los Alamos: The LANL Master Accounts Domain is Up and Running!"—available at this Web address: http://www.lanl.gov/orgs/cic/computingatlanl/ .

Creating New User Groups

Designed to look and act like the existing application called User Manager, the Domain Manager application allows administrators to create new user groups out of existing groups and/or individuals. Already listed in the LANL Domain Manager are "read-only groups," which are existing Laboratory organizational groups such as X-CI, CIC-4, TSA-4, and user-defined groups. The read-only groups will soon be automatically updated every night by running against the Laboratory's Employee Information System. Figure 1 is a screen capture of the LANL Domain Manager main window. In this example you can see the list of existing groups in the bottom window. To see the display of existing group members of the group selected in the group window, select File—Show Users.

An NT administrator can create or modify a user-defined group—a group created to serve a particular work team or project. For example, a new project might involve a team of all members of XXX-40 group plus 12 individuals from other Laboratory groups.[1] The administrator can use the File—New Global Group command to name the new group, locate each individual user and add him/her to the group, and select XXX-40 from the list of groups (Figs. 2 through 5). All individuals in XXX-40 become members of the new group.

This new user-defined group is then displayed in the LANL Domain Manager group window with all the organizational groups and other user-defined groups from which an administrator can choose.

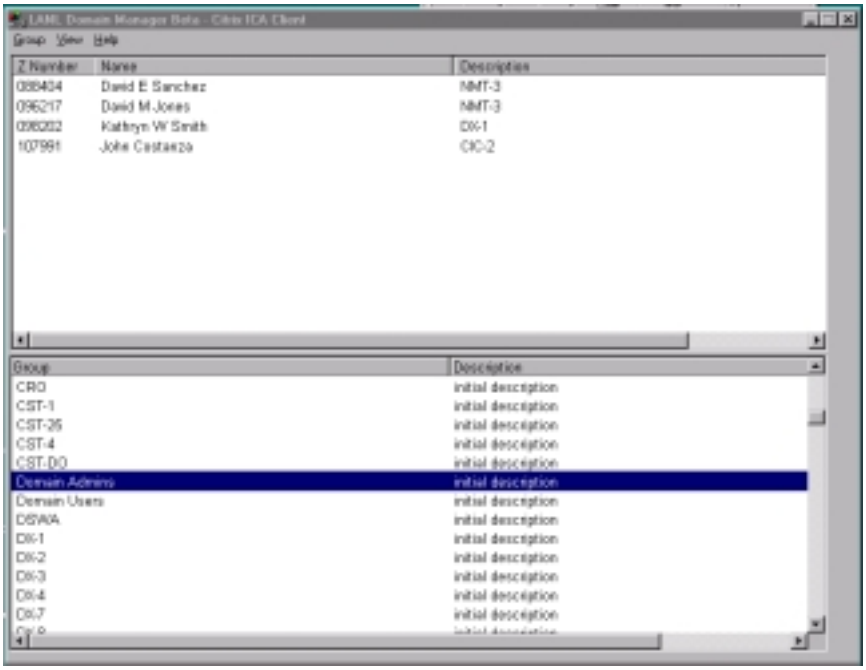

Fig. 1. The LANL Domain Manager main window. The list of existing groups appears in the bottom window; selecting File—Show Users displays the members of the selected group in the user (top) window.

Adding or Deleting Users in a Group

One of the most frequent tasks performed by network administrators is adding or deleting users in a group, and it's just as easy to do as creating a group. One method is to select the group name in the list on the LANL Domain Manager and from the File menu, select Properties. In the Global Group Properties dialog box, members of the group are listed along with Add and Remove buttons (shown in Fig. 4). You can select individual or multiple members or groups to remove at one time. The Add button allows you to look up users (shown in Fig. 3) you want to add; again, you can add individuals or multiple members or read-only groups. No user accounts are deleted from the MAD, but user memberships in groups are manipulated using the LANL Domain Manager.

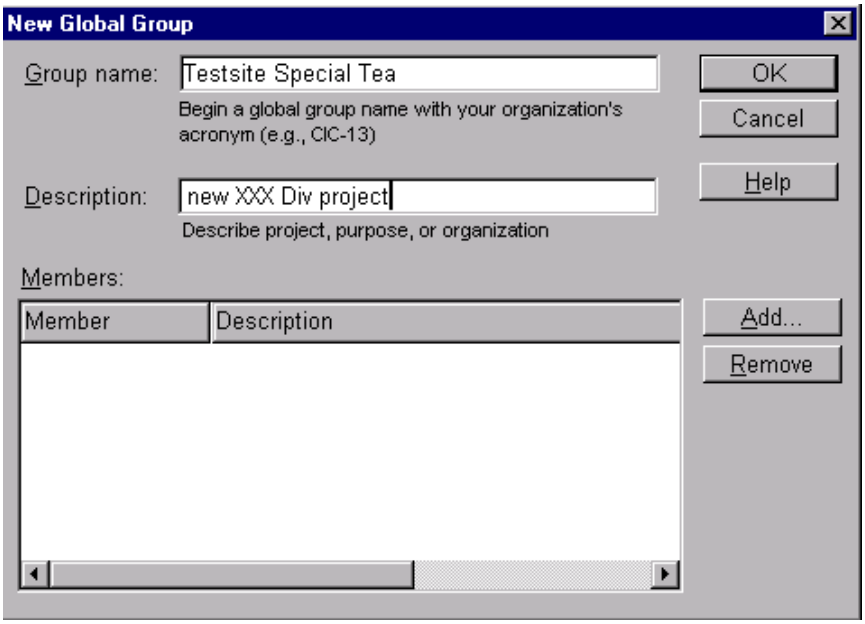

Fig. 2. The New Global Group dialog box.

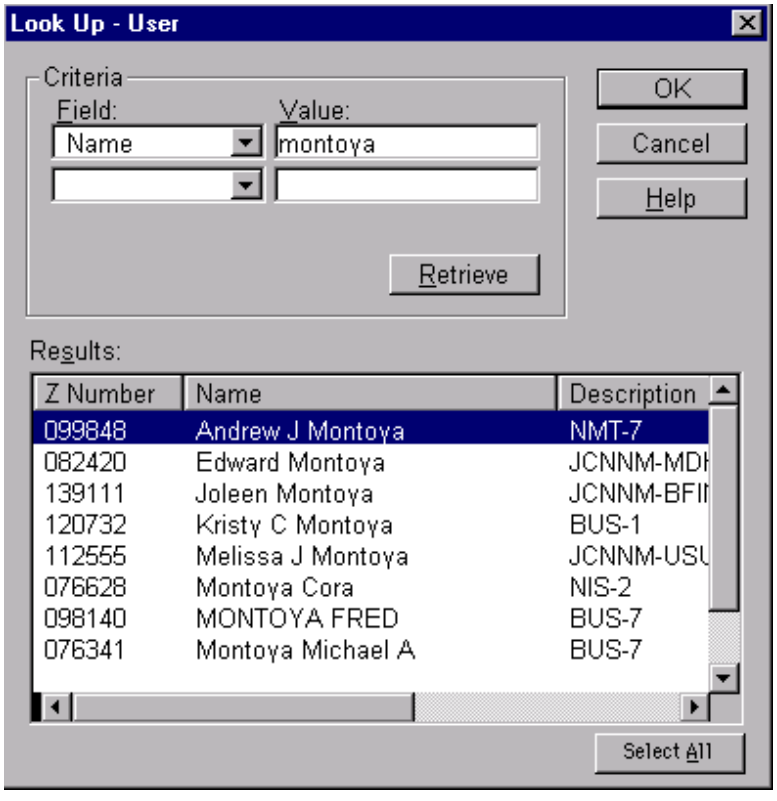

Fig. 3. You can look up users by name, Z number, or description and select one or more from the list to add to a new group.

Other capabilities available to administrators in the LANL Domain Manager include:

- looking up any user by name, Z number, or description
- copying any group membership and then adding and deleting members as needed as an easy way of creating a new group
- checking/modifying properties of one or more users
- checking user profiles
- listing all the groups a user belongs to
- deleting a user group (but not a global read-only group)
- generating and printing a report on one or several groups to see membership, broken down by individuals and member groups
- sorting user lists by name or Z number
- disabling or enabling user accounts

Online help is included under the Help menu to assist users in learning the new application. If you have any questions or suggestions for improving the LANL Domain Manager, contact Cheryl Host, NT MAD Project Leader, at chost@lanl.gov.

[1]Testsite Special Team was a group created only for illustrative purposes for this article. There is no such team, nor is there a group XXX-40.

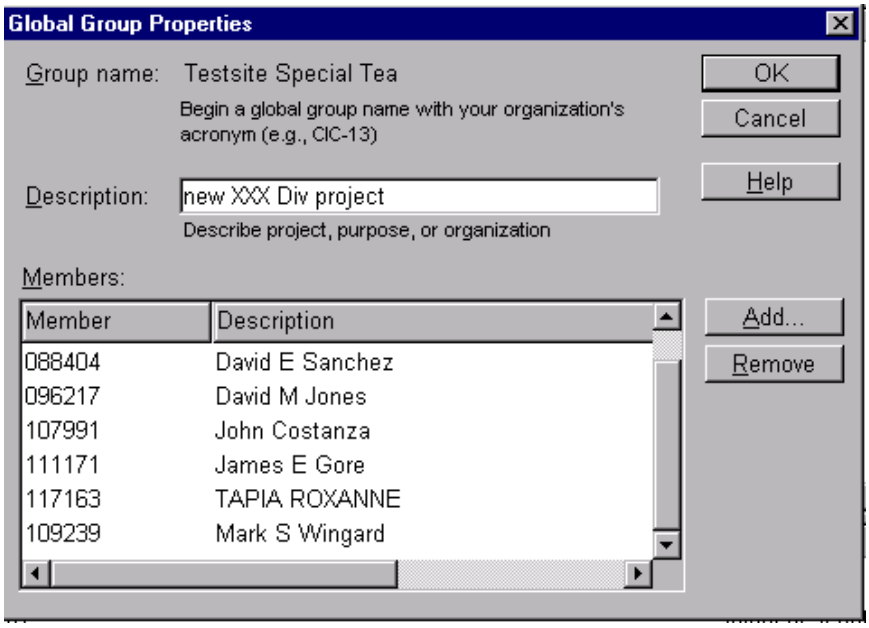

Fig. 4. The new group called Testsite Special Team.

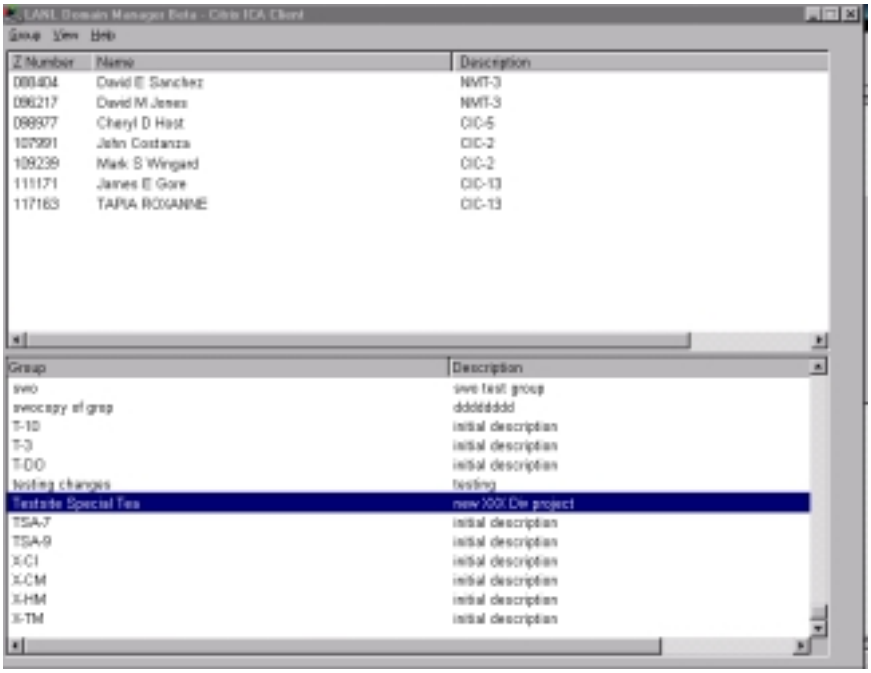

Fig. 5. The new group and its individual members (members of imaginary group XXX-40 are not displayed).

Keeping Informed about Progress on the Strategic Computing Complex Construction Project

by Denise Sessions, BITS Managing Editor, with Nicholas Nagy, SCC Deputy Project Director

The Strategic Computing Complex (SCC) is being built to support ASCI, the Advanced Strategic Computing Initiative. ASCI is a collaboration between three national laboratories and the Defense Programs arm of the Department of Energy. Its purpose is to create the leading-edge computations modeling and simulation capabilities that are essential for maintaining the safety, reliability, and performance of the US nuclear stockpile in the absence of nuclear testing (the Stockpile Stewardship Program) and reducing the nuclear danger.

The SCC will be a facility dedicated to simulating problems of extreme complexity and national urgency. Computational modeling of highly complex problems requires computers of immense power. These computers need space for the supporting infrastructure, including mechanical and electrical equipment, networks, workstations, visualization centers, interactive data-analysis tools, and collaborative laboratories.

The SCC should support any computer company's 30-TeraOps platform by January 2002, but the facility's design is flexible enough that it can house a 100-TeraOps computer by FY2004 and additional computing capability beyond. Mechanical and electrical equipment can be added to the building in increments to support a growing system. The facility's scalable nature includes future installation of chillers, cooling towers, computer room air-conditioning units, substations, motor generator power-conditions, transformers, and panel boards. Scalability is a cost-effective option that allows us to capitalize on technological advances in both computing technology and support equipment.

The SCC will provide space where multidisciplinary teams of weapons designers, computer scientists, physicists, code developers, engineers, and mathematicians can work together.

The SCC will be a modern facility and the cornerstone of the revitalized Technical Area- (TA-) 3 at Los Alamos National Laboratory. These are some of the features of the new building:

- 291,000 gross sq. ft.
- 43,500 sq. ft. unobstructed computer room
- Capable of air or water cooling of computers
- Scalable building utilities (support 30 to 150 TeraOps computer)
- Power — 7.1 MW expandable to 30 MW
- Water — 130,000 GPD expandable to 215,000 GPD
- Classified/Unclassified features
- Visualization theaters (2)
- Collaboratories (4)
- Design Simulation Laboratories (200 classified, 100 unclassified)
- 200-seat auditorium

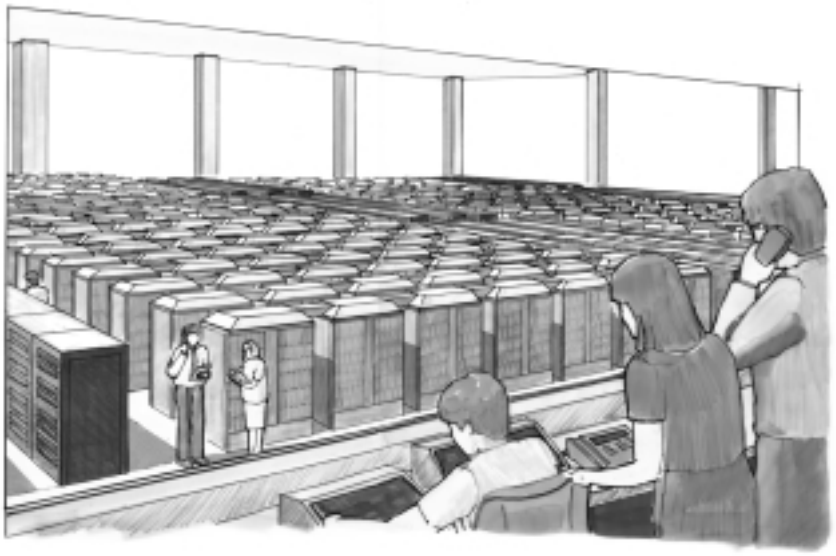

Artist's concept of computer room as viewed from the control room. Illustration by Design Workshop, Santa Fe, New Mexico, 1997

For More Information

The SCC Construction Project maintains a Web site that is a comprehensive source for getting the news on project progress as well as the project management documentation. See this URL: http://int.lanl.gov/projects/asci/scc/ . For example, you can track daily progress at the construction site via a Web-cam link. The Web site also includes links for news and photos, schedules, procurement, and documents. The documents link provides viewing of pdf files for related project planning such as environmental assessment, design-build request for proposal, and planning procedures. A link to the list of project team members shows the structure of the team and links for how to reach team members.

"TAKE A COUPLE OF DAYS OFF, RAY."

Policy Changes for Foreign Visits and Assignments

by Denise Sessions, BITS Managing Editor, CIC-1 Communications Arts & Services, with Lori Hutchins, ISEC/Foreign Visits and Assignments Team Leader

Laboratory personnel need to be aware of new policies regarding the access of foreign visitors and assignees to national laboratories. Recent changes in Department of Energy (DOE) policy and the signing of the US National Defense Authorization Act require changes to approval processes. Online training, new forms, and information on the subject are available on various Web sites.

DOE Notice 142.1

In July 1999, DOE Secretary Richardson issued directives to implement changes to the DOE policy and procedures concerning visits and assignments by foreign nationals to DOE facilities. In the memorandum announcing the policy changes, Secretary Richardson stressed that international cooperation and collaboration are important elements in the effective planning and implementation of many DOE programs. The new policy, DOE Notice 142.1, is online on the DOE Directive home page: http://www.explorer.doe.gov:1776/htmls/directives.html .

Process for Requesting Foreign National Participation

The DOE no longer approves visits or assignments from foreign countries. In accordance with DOE Notice 142.1, the Laboratory Director has the authority to approve foreign national visitors or assignees at LANL and at off-site DOE-sponsored events. LANL's Director John Browne has delegated this approval authority to the Associate Laboratory Directors (ALDs). As a result of this change, Form IA 473 is no longer valid and is replaced by Form 982, *Request for Unclassified Visit or Assignment by a Foreign National*. Pertinent safeguards and security forms can be accessed, printed, and/or filled in on the Laboratory's PDF online forms Web page. You can find this page on the LANL Administrative Resources Web site at: http://enterprise.lanl.gov/safe.htm .

Who Is Affected by This Policy?

The new policy applies to any visitor or assignee to LANL who does not hold US citizenship, including individuals who have obtained US residency but have not been naturalized.

When submitting their requests for foreign national participation, hosts should carefully consider the following criteria:

- Program needs,
- Export control issues,
- Technical merit,
- Potential sensitivities,
- Proximity of research or exposure to programs in support of weapons-related programs, and
- Unclassified computing vulnerabilities.

Approval and certification for the visit or assignment must be obtained before the visit or assignment start date. Table 1 outlines the details.

Approval Signatures for Form 982

The host must obtain signatures that certify line management support of the request. Signature authority cannot be delegated to anyone other than those individuals acting officially in those capacities. These signatures include

- Host,
- Cohost,
- Group leader,
- Division or program director, and
- Cognizant ALD or DLD.

Foreign Visits Office Approval Process

After approving your request, the Foreign Visits Office will send an electronic message (e-mail) to the host, cohost, and point-of-contact. This e-mail message is the host certification form. The host and cohost should read, sign, and return the signed (hard copy) certification form in inter-office mail to MS F675. For security reasons, the Foreign Visits Office must have hard copy, signed certifications instead of responses to e-mail messages. The point-of-contact does not need to sign the certification form.

List of Nonsecure Buildings

A Web site — http://s6server.lanl.gov/isec/newta.html — is updated as rooms and buildings become approved or removed from the list. Refer ONLY to this site for the most current listing. A red strike through the building or room will indicate if a building or room has been removed from the list. To request that a building be added to this list, complete Form 1767, *Request to Place Building on FV&A Exclusion List*, or download the form (PDF) from the Web site.

More Information Available for Hosts

An online training module designed for LANL staff who host or work with foreign nationals is available on the Web at the following URL: http://bus.lanl.gov/bus6/export-training/visit01.htm . This short (15-minute) self-study course briefs hosts on their responsibilities, the appropriate paperwork required, whom to contact with questions, badging requirements, and sensitive subject determination.

Moratorium on Visits and Assignments from Sensitive Country Foreign Nationals

In October 1999, the President of the US signed the National Defense Authorization Act for Fiscal Year 2000. This Act contains language relating to visits and assignments by "a citizen or agent of a nation that is named on the current DOE sensitive countries list." Beginning November 4, 1999, there is in effect a moratorium on visits and assignments by sensitive country foreign nationals to national laboratories. This moratorium will be in effect for at least 60 days or until the FBI, CIA, and DOE agree that the national laboratories have a valid protection measure in place to prevent foreign nationals from accessing classified information and from accessing sensitive unclassified information without a need to know and without proper authority.

Unclassified Open Computer Network Access

As of this writing, Internal Security (ISEC) is issuing new guidance for requesting access to the Laboratory's unclassified open (informally, "yellow") computer network. For the most updated information, see the Foreign Visits and Assignments Web site at http://s6server.lanl.gov/isec/fva.html .

Supercomputer Access

To access supercomputer or mass storage devices, hosts must obtain ALD Nuclear Weapons Technology approval. The request form is available in two formats (PDF and Microsoft® Word) for downloading on this Web site: http://c6help.lanl.gov/ under the "Dialing into LANL" header.

Additional Information

For information about whom to contact for questions, see the Foreign Visits and Assignments Web site at http://s6server.lanl.gov/isec/fva.html. For additional information, contact Lori Hutchins of ISEC at 5-5561 or lori@lanl.gov.

Table 1. Approval Process for Visits or Assignments of Foreign Nationals at DOE Facilities

Foreign National Visitor/Assignee	Requests Access to (Area):	Requests Access to (Information):	Process
Nonsensitive Country	Nonsecure Area	Nonsensitive Topic	Submit completed Form 982 at least **5 days** before visit/assignment start date.
DOE Sensitive Country	Nonsecure Area	Nonsensitive Topic	Submit completed Form 982 at least **6 weeks** before visit/assignment start date.
Nonsensitive Country	Secure Area	Unclassified	Submit completed Forms 982, 1726, and 1764 at least **6 weeks** before visit/assignment start date.
DOE Sensitive Country	Secure Area	Unclassified	Submit completed Forms 982, 1726, 1764, and an exception letter at least **6 weeks** before the visit/assignment start date.
Nonsensitive Country	Off-Site DOE-sponsored Events	Must be information in the public domain	Submit Form 982 at least **5 days** before the event. (Not required if the event is "open to the public.")
Sensitive Country	Off-Site DOE-sponsored Events	Must be information in the public domain	Submit Form 982 at least **6 weeks** before the event. (Not required if the event is "open to the public.")

Publications Release and Accountability at the Laboratory

by Jan Gallegos, Classification Administrator, S-7 Classification Group, Security & Safeguards Division

All scientific and technical information generated at the Laboratory and intended for public release must be reviewed and processed by the Classification Group (S-7) prior to publication or submission for publication. The Department of Energy (DOE) definition of this information category is sufficiently broad to encompass nearly all information generated at the Laboratory that is not purely administrative in nature. In addition, DOE requires that S-7 process all scientific and technical information "products," even if distribution to the public is not intended. Examples of the latter material would include classified or otherwise controlled reports. The publication release process includes review for classified information, Unclassified Controlled Nuclear Information (UCNI), export controlled information, other unclassified sensitive information, and information of patent interest. S-7 also assigns the official DOE subject category, an official Los Alamos publication number, and enters the publication into the Los Alamos Authors System database and S-7's microfilm archive. Documentation for the Authors System (AUS) on Enterprise Information Applications (EIA) is available at http://iosun.lanl.gov:2001/htmls/infoSys/icn/labwide/labwide.html. Finally, publications are submitted to the DOE Office of Scientific and Technical Information (OSTI) as required by DOE Order. Publications containing only technical information which falls under an approved Designated Unclassified Subject Area (DUSA) (see http://int.lanl.gov/orgs/s/s7/dusa.shtml) do not receive S-7 classification review, but must still be processed through S-7 prior to publishing. Technical information may be included in reports, articles, papers, books or book chapters, talks, slides, viewgraphs, handouts, and also documents to be published on the Internet.

As a reminder, DOE M 475.1-1 requires that *all* material generated at the Laboratory, that does not fall under an approved DUSA and is not purely administrative, must be reviewed for classification before dissemination. For publications, S-7 review accomplishes this requirement. For all other material, a local Authorized Derivative Classifier (ADC) must provide a review.

LANL has three main numbered publication series.

1. **LA-URs and LA-CPs** (Los Alamos unlimited releases and controlled publications)

   LA-UR (or CP) numbers are assigned to technical information which will be published or presented outside the Laboratory, published on the Internet, or retained for archival purposes. The numbers follow the pattern of LA-UR (or CP) –yy-xxxx, where yy denotes the year and xxxx are sequential numbers for that year. LA-UR numbers are assigned to publications that are to be unlimited in distribution. LA-UR numbers are also assigned to DUSA publications. LA-CP numbers are assigned to publications that carry some sort of distribution limitation because they are classified or otherwise sensitive in nature.

2. **LA-Series Reports**

   LA-Series Numbers (for example LA-xxxxx-MS) are given to formal technical reports that are published by the Laboratory. The series is numbered sequentially from its inception in 1943 and may include designators for preliminary, full, or progress reports.

3. **LALPs** (Los Alamos Laboratory Publications)

   LALP numbers are given to a wide variety of high-visibility, promotional public relations type documents distributed by the Laboratory. They are often intended for nontechnical audiences.

Special Cases

Web Pages

Web pages that are "publications" (i.e., they are self-contained, stand-alone documents that contain no imbedded links, etc.) that are intended for public release (posting on a "green" server) are typically processed through S-7 as hardcopy and assigned LA-UR numbers as with other publications. See the Information Architecture Standard "IA-6305: Classification Review and Publication Release for Internet/WWW Publications." If web pages are not stand-alone "publications," and contain links, etc., they are typically *not* assigned LA-UR numbers but are reviewed and approved online by S-7 following the procedures in the IA Standard "IA-7202: Online Classification Review Process for Laboratory WWW Materials." For nonpublication material posted on "yellow" (unclassified protected) servers, local Authorized Derivative Classifier (ADC) review is required, and sufficient.

Computer Codes

Public release of computer codes (software) is subject to special DOE requirements. Codes are assigned a LA-CC number by S-7. For more information on the process, refer to the Industrial Business Development (IBD) web pages at http://ext.lanl.gov:80/orgs/citpo/PDAT/IAM/software2.html.

Authors' organizations approve the technical content and propriety of publications (see http://int.lanl.gov/orgs/s/s7/pub_release.shtml). For processing of a Laboratory publication, the appropriate submittal form should have the name and signature of the person responsible for assuring that organizational procedures have been followed. *However, individual authors are responsible for obtaining the proper classification and sensitive information reviews of their documents and proper marking of their documents, as needed.*

For each publication, complete the Technical Information Release Form (Form 678, for LA-URs and LA-CPs), or other appropriate release form, and appropriate cover sheets, available at http://int.lanl.gov/orgs/s/s7/pub_release.shtml. Attach two copies of an unclassified abstract or viewgraphs intended for presentation, three copies of an unclassified paper, and one copy of a controlled or classified abstract, viewgraphs, or paper. For viewgraphs, include an abstract. Submit the complete package to the Classification Group, S-7, at TA-3, SM-43, A115, MS-F674. ***Please allow three working days for the classification review after the receipt by S-7.*** DUSA publications are given the LA-UR numbers in one day. Classification and sensitive information reviews are done by S-7, and the Business and Patent Law does the patent review. After the review, S-7 notifies the author of the publication number assigned, which constitutes approval for release (within the distribution limitation and with appropriate need-to-know for LA-CPs). For a complete listing of Laboratory approved document designators, see the above URLs. After reviewing the information on the S-7 homepage (http://int.lanl.gov/orgs/s/s7), please call the Classification Group at 667-5013 if you have further questions.

The above policy and procedures are based on the requirements of the following DOE directives and Laboratory policy:

a. DOE O 241.1 Scientific and Technical Information Management,

b. DOE M 475.1-1 Identifying Classified Information,

c. DOE O 471.1 Identification and Protection of Unclassified Controlled Nuclear Information.

d. DOE M 471.2-1B Classified Matter Protection and Control Manual

e. AM 721.08 "Conflict of Interest: Privileged Information – Release of Technical Data"

References

The Los Alamos National Laboratory publications release and accountability policy and procedures are available in online resources:

a. "Publishing at Los Alamos" at http://int.lanl.gov/publishing

b. "IA-6301: Guidelines for Publishing on the Laboratory Internet/WWW" at http://www.lanl.gov/projects/ia/stds/ia630112.html

c. "IA-6305: Classification Review and Publication Release for Internet/WWW Publications" at http://www.lanl.gov/projects/ia/stds/ia630513.html

d. "IA-7202: Online Classification Review Process for Laboratory WWW Materials" at http://www.lanl.gov/projects/ia/stds/ia720210.html

e. "Publications Release and Accountability" at http://int.lanl.gov/orgs/s/s7/pub_release.shtml

LOS ALAMOS NATIONAL LABORATORY

Research Library

http://lib-www.lanl.gov

The LANL Research Library offers a variety of training opportunities for the Laboratory community. Available sessions focus on specialized library databases and other electronic resources. A complete list of course offerings can be found at http://lib-www.lanl.gov/libinfo/training.htm. All sessions are available to individuals or groups at the library or your site. Arrange for a session by contacting the Library, phone 7-4175 or e-mail library@lanl.gov. Library tours are available on a drop-in basis every Wednesday at 1:00 p.m.

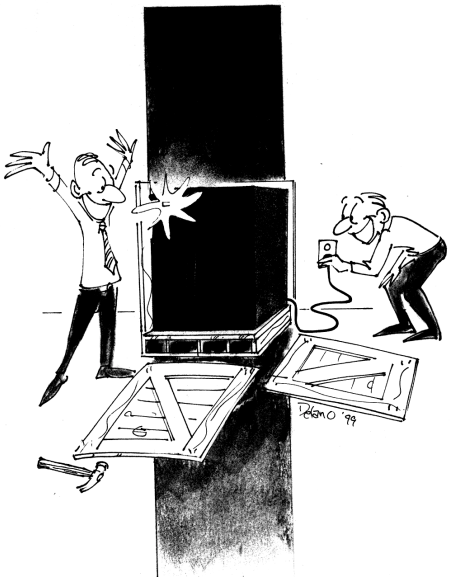

Attention Application Developers!!!

Need User Guides?

Need Online Help?

CIC-1's Software Documentation Team can help you improve your product by developing user manuals, online help, quick reference cards, or other documentation to help your users learn the application. Contact Sheila Molony, 5-1585, or sheila.molony@lanl.gov, for more information.

Computer Training

The Customer Service Group (CIC-6) offers technical computer training (enterprise information applications, communications, office administration, and Web authoring) and advanced computer training (programming languages, system administration, and advanced applications). To register for a course access our Web page at http://www.lanl.gov/internal/training/training.html . Or from the LANL home page select the links: Training, Computer. For further information about technical computer training call (505) 667-9559, and for advanced technical computer training call (505) 667-9399.

Technical and Advanced Technical Computer Training Courses

Communications	Office Skills 2000	Web Authoring and Browsing
• Eudora 4.2 • Meeting Maker 5.5.3	• Office Skills 2000–LANL Computing • Office Skills 2000–Professional Development	• Dreamweaver 2.0 • FrontPage 2000 • HTML
Enterprise Information Applications (EIA)	**Other EIA Courses**	**System Administration Training**
• Date Warehouse–Basics • Date Warehouse–EDS Reports • Data Warehouse–Passport • EDS–Basics • EDS–GUI • EDS–Training Plans • Infomaker • Invoice Approval System • Procurement Desktop • Purchase Card System • Recharge • Time & Effort GUI • Travel Foreign • Travel Domestic GUI • Web JIT	• Automated Chemical Information System (ACIS) • Directory Information System (DIS) • Financial Management Information System (FMIS) • Key/Core • Property Accounting, Inventory, and Reporting System (PAIRS) • Resource Planning Module (RPM) • Salary Management Tool (SMT) • Secretarial/Contract Service (SE) • Signature Authority System (SAS)	• IRIX (SGI) System Administration (Beginning) • IRIX (SGI) System Administration (Advanced) • IRIX (SGI) Network Administration • SGI Performance Evaluation and System Tuning • Solaris 7 System Administration I • Solaris 7 Network Administration II • Solaris TCP/IP Network Integration • Unix and Widows NT Integration • Windows NT Optimization and Troubleshooting • Windows NT Windows 2000 Update for IT Professionals • Windows NT Workstation and Server
Programming Training	**Application Training**	
• Advanced Perl Programming with CGI • ANSI/ISO C++ Programming Clinic (Advanced C++) • C Programming (Beginning) • C Programming (Advanced) • C++ for Experienced C Programmers • Distributed Objects Using Corba • Distributed Programming with Java • Java Programming • Java Program Workshop • Object-Oriented Analysis and Design • Object Technology–A Management Overview • Perl Programming • Programming for Beginners Using Java • Shell Programming	• Advanced WWW Development • Apache Web Server • C++ and the Unified Modeling Language • Foundations of IDL Programming • FrameMaker Basic and Advanced • IDL 5.0 Graphic Object Workshop • Netscape Servers for Intranet Development • Origin 2000 Applications Programming and Optimization • Sendmail–Managing Internet Mail • Sybase Fast Track to Adaptive Server Enterprise 11.5 (ASE) • Sybase Performance and Tuning for System 11 • Sybase SQL Server Administration • Unix (Beginning) • Unix (Advanced) • Visual Basic 5.0 Fundamentals • Visual C++ Windows Programming	

*You do not need an ICN password to use e-mail.

Reader Feedback

Feedback helps us to provide a document that responds to the changing needs of its readership. If you have comments or questions about this publication, please let us hear from you. We have reserved the back of this form for that purpose. We also accept articles for publication that are of interest to our readers. Contact the managing editor for more information. This form is also used for new subscriptions, deletions, or changes. Instructions are on the back. If you prefer to contact us by e-mail, send your comments and/or subscription request to denise@lanl.gov.

Do Not Staple
Fold on This Line First

cut along dashed line

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST-CLASS MAIL PERMIT NO. 88 LOS ALAMOS NM

POSTAGE WILL BE PAID BY THE ADDRESSEE

MAIL STOP B251
ATTN: DENISE SESSIONS, MANAGING EDITOR
CUSTOMER SERVICE GROUP (CIC-6)
LOS ALAMOS NATIONAL LABORATORY
PO BOX 1663
LOS ALAMOS, NM 87545-9916

Do Not Staple, Seal with Tape
Fold Here

1999 12-Month Index

To electronically access any of the articles listed on this index go to http://www.lanl.gov/orgs/cic/cic6/bits/index/index_home.html. Note that the page numbers printed on this page correspond to the page numbers in the printed versions of the publication.

Keywords	Title of BITS Article	Date	Page
High Performance Computing	*Performance and Scalability Analysis of Applications on T-flop-scale Architectures*	*J/J 99*	*7*
	Major 3-D Parallel Simulations	*J/J 99*	*9*
	Complex Systems Modeling: Using Metaphors from Nature in Simulation and Scientific Models		
	Dimensionless Coding Techniques in Tecolote: Using a Single Source-Code Base for Multidimension Programs	*Nov. 99*	*22*
	Performance and Scalability Analysis of Applications on T-flop-scale Architectures	*J/J 99*	*7*
	Major 3-D Parallel Simulations	*J/J 99*	*9*
	Lecture Review: Next-Generation Chips, Processors, Transistors, and Wiring	*Aug. 99*	*23*
	Tecolote: An Object-Oriented Framework for Hydrodynamics Physics	*Aug. 99*	*26*
	Frameworks Are Models, Too!	*Oct. 99*	*16*
	Automated OS Install HOWTO: Linux Cluster via Gigabit Ethernet	*Oct. 99*	*20*
Infrastructure	*Virtual Private Network*	*J/J 99*	*15*
	Electronic Information Protection Regimes	*Aug. 99*	*13*
	Web Content Accessibility: New W3C Guidelines Have Arrived	*Aug. 99*	*20*
	Web Content Architecture: Taming the Tangle of Protocols	*Sept. 99*	*15*
	Data Analysis in the Advanced Surveillance Technology Initiative: AMISS Testbed	*Oct. 99*	*9*
	Transaction Integrity and Validity in MC&A Databases	*Oct. 99*	*12*
	Take a Look at the New PAGES	*Nov. 99*	*11*
	Recharge System Telecommunications Enhancements	*Nov. 99*	*13*
Information Systems	*Engineering Index Is Now Available from the Research Library*	*Dec. 98*	*13*
	Research Library Adds Social SciSearch® Database	*F/M 99*	*5*
	Library Databases	*J/J 99*	*20*
	aha!—An Interface to the Best of LANL Science & Technology Web Resources	*Aug. 99*	*2*
	More Than a New Look & Feel: The New LANL Badging System	*Aug. 99*	*5*
	Laboratory Overview Presentation Materials on the Web	*Sept. 99*	*2*
	New INSPEC® at LANL	*Sept. 99*	*28*
	Keep on Top of Information Technology with Gartner Group	*Sept. 99*	*28*
	Beilstein Organic Chemistry Database Is Now on the Web—Free	*Nov. 99*	*27*
	New Interface for Biosciences Database	*Nov. 99*	*27*
Networks	*Laboratory Unclassified Network Will Implement Changes for Security*	*Dec. 98*	*3*
	Preserving Bandwidth on Your Windows NT Network	*F/M 99*	*8*
	News About Networking at Los Alamos: The LANL Master Accounts Domain Is Up and Running!	*Sept. 99*	*10*
	Notes on Setting Up and Configuring an NT Domain	*Sept. 99*	*12*

Customer Support Center . . . . (505) 665-4444, ext. 851, or cichelp@lanl.gov

Because of a wide variety of CIC computing services, numerous facilities are available to address your questions. If you are uncertain whom to call, you can always call the Customer Support Center (CSC). CSC consultants are trained to either answer your question or locate someone who can. To reach the appropriate consultant, dial 665-4444 and make your selection from the following choices:

Option 1: New user topics including e-mail, passwords, registration, and World Wide Web
Option 2: Enterprise Information Applications such as Travel, Time and Effort, and Purchase cards
Option 3: Scientific computing, storage systems, and networking
Option 4: Classroom instruction and training
Option 5: Desktop Consulting for PC and Macintosh software and network configurations.

Consulting Via E-Mail

Customer Support ..... cichelp@lanl.gov
Scientific and engineering computing ..... consult@lanl.gov
Administrative and business computing .... eiaconsult@lanl.gov
Passwords and registration ..... validate@lanl.gov
Macintosh computing ..... Mac-help@lanl.gov
PC computing ..... PC-help@lanl.gov
UNIX computing ..... UNIX-help@lanl.gov

Other Useful Numbers

Advanced Computing Laboratory 665-4530
Central Computing Facility 667-4584
Network Operations Center ..... noc@lanl.gov or 667-7423
Telephone Services Center 667-3400

Los Alamos
NATIONAL LABORATORY

STRATEGIC COMPUTING COMPLEX

A facility dedicated to simulating problems of extreme complexity and national urgency.

The **Strategic Computing Complex** will meet the urgent national security requirements associated with the Stockpile Stewardship Program. It will house an immense, integrated system of computer processors.

Collaboratories will provide users and managers with an information- and technology-rich environment for simulation developers, collaboration, discussion, media-development, and problem analysis.

Visualization Theatres will support the latest virtual-reality technologies for three-dimensional data presentation, interaction, and exploration.

Simulation Laboratories will provide designers, analysts, and code developers the capability to interact with the immense data sets that will be produced by simulations on TeraOps platforms.

For more information about the construction project, see the article inside this issue.

Los Alamos
NATIONAL LABORATORY

Mailstop B251 Los Alamos, New Mexico 87545

A U.S. Department of Energy Laboratory
BITS is published bimonthly to highlight recent computing and communications activities within the Laboratory. We welcome your suggestions and contributions.

BITS may be accessed electronically at this URL:
http://www.lanl.gov/cic/publications.html

LALP-00-2 (1/00)

Nonprofit
organization
US Postage
PAID
Los Alamos, NM
Permit No. 532